On the Statistical Analysis of X-ray Polarization Measurements

T. E. Strohmayer and T. R. Kallman

X-ray Astrophysics Lab, Astrophysics Science Division, NASA's Goddard Space Flight Center, Greenbelt, MD 20771

Received _____; accepted _____

ABSTRACT

In many polarimetry applications, including observations in the X-ray band, the measurement of a polarization signal can be reduced to the detection and quantification of a deviation from uniformity of a distribution of measured angles of the form $A + B\cos^2(\phi - \phi_0)$ $(0 < \phi < \pi)$. We explore the statistics of such polarization measurements using both Monte Carlo simulations as well as analytic calculations based on the appropriate probability distributions. We derive relations for the number of counts required to reach a given detection level (parameterized by β the "number of $\sigma's$" of the measurement) appropriate for measuring the modulation amplitude a by itself (single interesting parameter case) or jointly with the position angle ϕ (two interesting parameters case). We show that for the former case when the intrinsic amplitude is equal to the well known minimum detectable polarization (MDP) it is, on average, detected at the 3σ level. For the latter case, when one requires a joint measurement at the same confidence level, then more counts are needed, by a factor of ≈ 2.2, than that required to achieve the MDP level. We find that the position angle uncertainty at 1σ confidence is well described by the relation $\sigma_\phi = 28.5(\text{degrees})/\beta$.

Subject headings: polarimetry: general — statistical analysis: Monte Carlo simulations

1. Introduction

Emission and scattering processes thought to be important in many astrophysical X-ray sources are likely to impart specific polarization signatures, but to date there have only been a few positive detections of polarization from cosmic X-ray sources, largely due to sensitivity limitations. Some of the earliest and highest precision measurements were made with the OSO-8 Bragg reflection polarimeter (Kestenbaum et al. 1976; Weisskopf et al. 1976), and include high significance measurements of the linear polarization properties of the Crab nebula in several energy bands (Weisskopf et al. 1978).

More recently, observations with the *INTEGRAL* spectrometer SPI and imager IBIS have exploited the polarization dependence of Compton scattering to infer the linear polarization properties of the Crab at γ-ray energies (Dean et al. 2008; Forot et al. 2009). These results indicate that the > 200 keV flux from the Crab nebula is highly polarized ($\approx 50\%$), with a position angle consistent with the pulsar rotation axis.

In the last few years the development of micropattern gas detectors has enabled the capability to directly image the charge tracks produced by photoelectrons, thus enabling use of photoelectric absorption in a detection gas as a direct probe of X-ray polarization (Costa et al. 2001; Black et al. 2004; 2010). It is likely that such technology will be employed in the not-too-distant future to sensitively explore the polarization properties of many classes of astrophysical X-ray sources for the first time.

In this paper we explore the question of how one detects, measures, and characterizes a polarization signal with a photoelectric polarimeter. The remainder of this paper is organized as follows; in §2 we outline the basic problem of detecting a modulation in a distribution of angles, and we describe the angle distributions used throughout the paper. In §3 we briefly outline the probability distribution relevant to such polarization measurements. In §4 we describe our Monte Carlo simulations and present our results. We

conclude with a short summary in §5.

2. Statement of the Problem

The angular distribution of photoelectrons ejected by a linearly polarized beam of photons (X-rays) is proportional to $\sin^2(\theta)\cos^2(\phi)/(1-\beta\cos(\theta))^4$ (see, for example, Costa et al. 2001), where θ is the emission angle measured from the direction of the incident photon ($0 < \theta < \pi$), and ϕ is the azimuthal angle measured relative to the polarization vector of the incident photon (see Figure 1 for the basic geometry applicable to a photoelectric polarimeter). In most practical situations the angle θ is not inferred directly, but the electron charge track projected into the plane orthogonal to the direction of the photon (the plane defined by $\theta = 90$ deg) is imaged and thus ϕ can be estimated for each detected photon. The angle ϕ is measured around the line of sight to the target of interest and can be referenced to, for example, local North on the sky. The presence of a significant linear polarization component is then evident in the distribution of azimuthal angles. For example, an unpolarized photon flux will produce a uniform distribution in the angle ϕ, whereas a linear polarization component produces a distribution peaked at a particular azimuthal angle, ϕ_0.

In general, the observed distribution in ϕ will be of the form,

$$S(\phi) = A + B\cos^2(\phi - \phi_0)\,. \tag{1}$$

This distribution has an amplitude of modulation given by, $a \equiv (S_{max} - S_{min})/(S_{max} + S_{min}) = B/(2A + B)$, and a position angle (the angle at which the distribution has a maximum) given by ϕ_0, thus, the detection of polarization can be reduced to a statistical detection of a modulation in the distribution of angles, ϕ. Such a distribution is often referred to as a modulation curve.

In principle the angle ϕ can be measured in the range from $0 - 2\pi$ (0 - 360 degrees), however, due to the two-fold symmetry of the $\cos^2(\phi)$ dependence of the angular distribution of ejected photoelectrons, it is sufficient to consider distributions over a range of angles from 0 to π (0 - 180 degrees).

An equivalent and often convenient way to express this distribution is using the so-called Stokes decomposition,

$$S(\phi) = I + Q\sin(2\phi) + U\cos(2\phi)\,, \tag{2}$$

where I, Q and U are the well-known Stokes parameters, the modulation is now given by $a = (Q^2 + U^2)^{1/2}/I$, and the position angle is $\phi_0 = 1/2\tan^{-1}(Q/U)$.

Now, the amplitude of modulation a is not in general equal to the source polarization amplitude, a_p, because a detector is not a perfect analyzer and will not provide an exact measurement of the true photoelectron angle ϕ. That is, individual angle estimates will have some uncertainty associated with them and these will produce a uniform (unmodulated) component to the measured distribution even in the case of a 100% polarized beam. This "lossiness" of the angle estimates is quantified in terms of the so-called detector modulation, μ, which is the modulation amplitude produced in the detector by a 100% polarized X-ray beam.

In general, μ can depend on a number of factors, including the energy of the incident photons, and the composition and pressure of the absorbing gas, among others (see Paciani et al. 2003, Bellazini et al. 2003). In the absence of background, the amplitude of polarization is just $a_p = a/\mu$. In general, a_p is larger than the measured amplitude of modulation, a, because as noted above, detectors are not 100% efficient, and some of the intrinsic polarization amplitude is smeared out.

3. Probability Distribution

Previous studies have shown that the joint probability distribution for a measurement of polarization amplitude, a, and position angle, ϕ is given by

$$P(N, a, a_0, \phi, \phi_0) = \frac{Na}{4\pi} \exp\left(-\frac{N}{4}(a^2 + a_0^2 - 2aa_0 \cos 2((\phi - \phi_0)))\right), \tag{3}$$

where a, a_0, ϕ, ϕ_0, and N are the measured amplitude, the true amplitude, the measured position angle, the true position angle, and the number of detected photons, respectively (see, for example, Weisskopf, Elsner & O'Dell 2010). In the case of no intrinsic polarization, $a_0 = 0$, the distribution simplifies substantially to

$$P(N, a) = \frac{Na}{4\pi} \exp\left(-\frac{N}{4}a^2\right), \tag{4}$$

and this expression can be readily integrated to find the probability of measuring an amplitude, a, if there is no intrinsic polarization. The amplitude that has a 1% chance of being measured is referred to as the *minimum detectable amplitude* (MDA), and it is relatively straightforward to show that $MDA = 4.29/\sqrt{N}$. The polarization amplitude that would produce this modulation amplitude in a particular detector system is called the *minimum detectable polarization* (MDP), and based on the discussion above is just $MDP = MDA/\mu = 4.29/(\mu\sqrt{N})$.

Furthermore, if we are not concerned with the position angle, ϕ, then we can integrate equation (3) over angles and obtain the distribution,

$$P(N, a, a_0) = \int_{-\pi/2}^{\pi/2} P(a, \phi) d\phi = \frac{a}{\sigma^2} exp(-\frac{(a^2 + a_0^2)}{2\sigma^2}) I_0(\frac{aa_0}{\sigma^2}) \tag{5}$$

where I_0 is the modified Bessel function of order zero, and $\sigma^2 = 2/N$. This distribution is known as the Rice distribution (Rice 1945) and it reflects the fact that the amplitude is always a positive quantity, and so the distribution must go to zero for $a = 0$. An important property of this distribution concerns the second moment, which is related to the width,

and is given by: $< a^2 >= a_0^2 + 4/N$ which shows that the distribution width increases with a_0.

4. Monte Carlo Simulations

Without loss of generality an X-ray source's linear polarization characteristics can be described by the modulation amplitude, a_0, and position angle, ϕ_0 (in the range from 0 - π), that would be produced in a particular detector system. In making an observation a total of N photons are observed for each of which an angle ϕ_i is estimated for the ejected photoelectron in the range from $0 - \pi$ (0 - 180 deg). One can then create a histogram of the number of events (photons) in each of M position angle bins, where the bin size (in degrees) is $\Delta\phi = 180/M$. One can then do least-squares (χ^2) fitting to estimate both a, ϕ, and their 1σ uncertainties (in practice we fit the modulation curves with the Stokes form of the distribution, equation 2). We call this a measurement of polarization. In effect, what is measured is the modulation amplitude, a, and it is the knowledge of the detector system, expressed in terms of μ, that enables this to be converted to a source polarization amplitude via the expression $a_p = a/\mu$.

This procedure is quite amenable to simulation with Monte Carlo techniques, and here we present results of such simulations, both in the case with $a_0 = 0$ and $a_0 > 0$. For now we ignore background considerations and work only in terms of modulation amplitudes, that is, the following results can be considered applicable to any detector system, regardless of the μ value that characterizes its polarization sensitivity.

The simulations proceed as follows, for a given set of true parameter values, a_0 and ϕ_0, we compute a number, M_{sim}, of simulated data sets each of which has N photons. For each simulated data set we bin the resulting angles to form a modulation curve and for each the

Stokes form of the distribution (equation 2) is fitted to determine best-fit values for Q, U, and I, and their 1σ uncertainties, σ_Q, σ_U, and σ_I, respectively. We can use the fitted Stokes parameters to express the results in terms of a and ϕ using the expressions defined above. The 1σ uncertainties, σ_a and σ_ϕ, can also be determined by standard error propagation methods, which yields,

$$\sigma_a = a\left((\sigma_I/I)^2 + (\sigma_Q^2/(Q^2+U^2))\right)^{1/2}, \tag{6}$$

and

$$\sigma_\phi\ (\mathrm{deg}) = (180/\pi)\left(\sqrt{(}x^2)/(2(1+x^2))\right)\left((\sigma_Q/Q)^2 + (\sigma_U/U)^2\right)^{1/2}, \tag{7}$$

where, $x = Q/U$, and σ_I, σ_Q, and σ_U are the standard 1σ uncertainties on the fitted quantities, I, Q and U. The procedure can then be repeated with different values of N. All the Monte Carlo simulations described here were done using *IDL*, and uniform deviates were obtained with *IDL's* random number generator *randomu*. We use a least-squares fitting routine within *IDL* developed by Craig Markwardt that is based on MINPACK-1 (see http://cow.physics.wisc.edu/čraigm/idl/fitting.html). In all of the least squares fits all parameters are allowed to vary.

4.1. Results with $a_0 = 0$: Minimum Detectable Amplitude (MDA)

In the case of $a_0 = 0$ the distribution of angles, ϕ, is uniform, which can be easily simulated with a random number generator that produces uniform deviates. We obtain from this procedure a distribution of best-fit amplitudes, a, from which we can estimate the value $a_{1\%}$ that has a 1% chance probability of being measured. This is just the familiar MDA described above. Figure 2 shows a comparison of the results from the simulations (blue square symbols) versus the analytic expression given above, $4.29/\sqrt{N}$ (solid line). We computed $M_{sim} = 10,000$ simulations for these results, and we used 16 position angle

bins for the modulation curves. Figure 2 shows that the $a_0 = 0$ simulations are in good agreement with the analytic result, giving us confidence that our simulation procedures are correct.

4.2. Results with $a_0 > 0$: Detection of Polarization

We can now explore a number of issues with regard to detection of polarization. For example, how many counts are needed to measure a modulation amplitude to a particular precision, and for a given precision in the amplitude measurement, how accurately can the position angle be measured?

To address these questions we perform additional simulations using true distributions with specified amplitudes and position angles. For the illustrative examples below we used two different amplitudes, both with $\phi_0 = 0$, however, we have explored many different cases and all the results summarized here are independent of the particular values of a_0 or ϕ_0 used. The example distributions described below are; $S(\phi) = 10 + 1\cos^2(\phi - 0)$ and $S(\phi) = 10 + 2\cos^2(\phi - 0)$. These examples correspond to intrinsic modulation amplitudes, a_0, of 1/21 and 2/22, respectively. For a specific detector these would correspond to polarization amplitudes, $a_p = 1/(21\mu)$, and $2/(22\mu)$.

To determine the number of counts, N, needed to reach a polarization sensitivity given by the MDP value we require that MDP $= 4.29/\mu\sqrt{N} = a_p = a/\mu$, and thus $a = 4.29/\sqrt{N}$. Note that here μ cancels out, and to achieve MDP values at these amplitudes requires $(21^2) * 4.29^2 = 8116.21$, and $(22/2)^2 * (4.29)^2 = 2226.90$ counts, respectively. Note that this may seem a trivial point, but it's important to keep the terminology as precise as possible. Below we will compare the counts needed to reach a certain MDP, N_{mdp}, and the counts needed to measure the same amplitude of polarization to a certain precision, N. While N

and N_{mdp} will individually depend on μ, the ratio, N/N_{mdp}, cannot depend on μ, that is, it is independent of the detector system employed.

When $a_0 > 0$ we sample random angles from the true distributions by the so-called transformation method. We first compute the cumulative distribution of $S(\phi)$, then draw uniform deviates, x, and find the corresponding value of $\phi(x)$ in the cumulative distribution. We use the root finder *ZBRENT* implemented in *IDL* to solve for ϕ given x. We can thus draw a specific number of events, N, from the true distribution. For each value of N we compute a large number M_{sim} of realizations, to each of which is fitted the Stokes distribution to determine the best fit values of I, Q, and U (and thus a and ϕ), and their 1σ uncertainties. We can thus simulate the distributions of a and ϕ for different N.

4.3. Single Parameter Confidence Regions

Here we present the results of our simulations in several ways. First, we used the results from $M_{sim} = 1,000$ realizations for different values of N to find the mean values of a, and ϕ and their 68% confidence ranges. These are derived for each parameter independently of the other, that is, they are 1-dimensional (single parameter) ranges. For example, if one were interested in asking the question, is a source (or population of sources) polarized, without regard to the particular position angle of the electric vector, then the 1-d distribution would be appropriate. We explore the joint, 2-dimensional distributions below in §4.4.

Figure 3 shows an example simulated distribution of measured amplitudes, a, computed with $N = 20,000$ events and $a_0 = 2/22$ (with $\phi_0 = 0$). The left panel shows the differential (binned) distribution, and the right panel shows the same results expressed as a cumulative distribution. We also computed similar distributions for the measured position angle, ϕ.

Procedurally we use the estimated cumulative distributions to identify the mean of

each distribution (ie., for both a and ϕ), and the two values a_{max} and a_{min} that enclose 68% of the distribution. This then gives the 1σ uncertainty, $\sigma_{a,1d} = (a_{max} - a_{min})/2$ (we also compute the corresponding values for the ϕ distribution). We can then compute the quantity $\beta_{1d} = a_{mean}/\sigma_{a,1d}$, which can be thought of as the "number of sigmas" of the measurement. We now explore the behavior of several quantities as a function of β_{1d}.

The first is the ratio N/N_{mdp} (see Figure 4), where N is simply the number of events (photons) simulated (ie. the number of observed counts in the modulation curve), and $N_{mdp} = 4.29^2/a_0^2$ is just the number of counts that would be required to reach an MDA equal to the true amplitude, a_0. This ratio can be thought of as the additional observing time required to measure the true amplitude to a given significance compared to the time needed to reach an MDA equivalent to the true amplitude. The black diamond symbols in Figure 4 show the results of simulations using the example distributions described above with $M_{sim} = 1,000$ for different values of N.

One can also present the results in a slightly different but complementary way. In the above simulations we have essentially carried out many simulated observations and computed the distribution of "observed" values of the amplitude and position angle, but in reality an observer may not have the luxury of making such a large number of independent observations of a particular source. All any observer can do is to observe some number, N, of photons from the source, construct a modulation curve and fit it as we have done in the simulations described above. From this procedure we obtain four quantities, a, σ_a, ϕ, and σ_ϕ. This constitutes a measurement of the polarization parameters, or more simply, a measurement of polarization. Here, the uncertainties, σ_a, and σ_ϕ, are obtained for a specified confidence level and number of degrees of freedom. For example, for the 1σ (68.3%), single parameter confidence ranges we would find the change in each parameter that produces a $\Delta\chi^2 = 1$ (while allowing the other parameter to vary in the fitting procedure). For a

2-dimensional, joint confidence region at the same level of confidence (68.3%) we would find the $\Delta\chi^2 = 2.3$ contour in the a - ϕ plane.

For the diamond symbols in Figure 4 above we used the mean values and 1σ uncertainties derived from the distributions computed from many simulated observations to obtain β_{1d}, however, one can also use the "measured" quantities from each simulated observation to compute $\beta_{obs,1d} = a/\sigma_{a,1d}$. One can then plot the observed quantities for each simulated observation, where now $N_{mdp}(a)$ is computed using the best-fit value for a and the formula $a = 4.29/N_{mdp}^{1/2}$. We have done this and show the results in Figure 4 with the colored symbols. That is, the colored points are these "measured" values from individual simulations. The red symbols were computed with $N = 10,000$, $a_0 = 2/22$, $\phi_0 = 0$ (deg), the blue used $N = 20,000$, $a_0 = 2/22$, $\phi_0 = 0$ (deg), and the green with $N = 24,000$, $a_0 = 3/25$, and $\phi_0 = 22.5$ (deg). We see that the distribution of "measured" points falls along the same relations as that deduced from the simulated distributions of a, and ψ, as indeed they should since they are sampling the same distributions, this way of presenting the results of the simulations makes a more direct connection with actual polarimetric observations, as we only plot quantities that one would obtain directly from a single observation. To the extent that actual observations are dominated only by poisson counting noise and for which the background is small, then they must fall along the relations followed by the simulated observations shown in Figure 4. Indeed, the locus of points traced out by the "measured" values from individual simulations can be easily approximated by simply plotting curves that intersect the entire swarm of simulated points. Doing this we find that $N/N_{mdp} = \beta_{obs,1d}^2/9.18$ (dashed curve in Figure 4) to very good accuracy.

We also explored simulations where the true amplitude a_0 was set equal to the MDA. Figure 5 shows the cumulative distribution of the "measured" values of β from several such simulations. One can see from Figure 5 that roughly 60% of the time one would "measure"

the amplitude, a, at the 3σ level, or better. We emphasize that this relation is based on the 1-d (single parameter) confidence range for a, and would be appropriate only for the case of addressing the question of the detection of a significant modulation amplitude independent of the position angle. We now explore the joint, 2-dimensional confidence regions.

4.4. Joint Confidence Regions for a and ϕ

For a 2-dimensional confidence region we need to find the contour in the a - ϕ plane that encloses a specified fraction of the best-fit pairs. For a 1σ region (68.3% confidence) this is the contour that satisfies $\Delta\chi^2(a,\phi) = 2.3$, where $\Delta\chi^2 = \chi^2(a,\phi) - \chi^2_{min}(a_{best},\phi_{best})$. We again use the Stokes decomposition and compute $\Delta\chi^2(I,Q,U)$ on grids of I, Q and U around the best-fit values, I_{best}, Q_{best}, and U_{best}. We can then convert the grids of Stokes parameters into the appropriate values of a and ϕ and find the boundaries of the region that satisfies $\Delta\chi^2 \leq 2.3$. We next find the maximum and minimum value on the boundary for each parameter. We can then define $\sigma_{a,2d} = (a_{max} - a_{min})/2$ and $\sigma_\phi = (\phi_{max} - \phi_{min})/2$, where a_{max}, a_{min}, ϕ_{max} and ϕ_{min} are the maximum and minimum values on the contour of a and ϕ, respectively. Figure 6 shows a pair of 1σ confidence regions computed in this fashion. Results from two simulations are shown, one with a lower amplitude, $a_{0,low} = 2/24$, and one with a higher amplitude, $a_{0,hi} = 3/25$. Both simulations were performed with $\phi_0 = 25$ deg (these true parameters are marked by the red diamond symbols), and N = 10000. In each case a modulation curve was randomly sampled using the true amplitude and position angle, and the simulated data were then fitted to determine the best-fit values of the amplitude and position angle. These points are shown by the green square symbols. The shaded areas show the regions of a and ϕ around each best-fit pair that satisfy $\Delta\chi^2 \leq 2.3$. The horizontal and vertical dotted lines mark the maximum and minimum values on the regions for each parameter. One can see that the size of the confidence region grows for

smaller intrinsic modulation amplitudes (as one would expect), and that the confidence regions are in general not circular.

The probability distribution described earlier (§3) can also be used to derive an analytic expression for the 2-d confidence contour in the a - ϕ plane for any desired level of confidence. This approach has been investigated by Weisskopf et al. (2010). They derive a pair of parametric relations for the values of a and ϕ on any confidence contour (see their equation 8). These expressions predict the correct range (extremes) in the amplitude, a, but overpredict the range in ϕ by a factor of 2, apparently because their original derivation neglected a switch from phase angle to position angle (Weisskopf, private communication). Thus, if one replaces all occurrences of ϕ, ϕ_0 and ψ with 2x the respective angle (eg., $\sin\phi_0 - - > \sin 2\phi_0$) beginning at their equation (6), one obtains the following expressions;

$$a = \left(a_0^2 + \Delta a_C^2 + 2a_0 \Delta a_C \cos 2(\psi - \phi_0)\right)^{1/2} \tag{8}$$

and

$$\phi = 1/2 \tan^{-1}\left((a_0 \sin 2\phi_0 + \Delta a_C \sin 2\psi)/(a_0 \cos 2\phi_0 + \Delta a_C \cos 2\psi)\right), \tag{9}$$

where $\Delta a_C = (-4\ln(1-C)/N)^{1/2}$, with C and N being the desired confidence level and number of detected photons, respectively, and ψ is just a parametric angle that varies around the contour. The thick blue curves in Figure 5 were drawn using these expressions with $N = 10,000$, $C = 0.683$ (ie., 1σ), and the pair (a_0, ϕ_0) given by the appropriate best fit values (green square symbols). These contour curves provide an excellent match to the boundaries of the shaded regions, indicating that the Monte Carlo simulations and analytic calculations are in excellent agreement.

We can now compute the value $\beta_{2d} = a_{best}/\sigma_{a,2d}$ for each particular simulation. This again quantifies the "number of sigmas" of the measurement, but now reflects the fact that it is a *joint* measurement of both a and ϕ together. Results from a number of such simulations are also shown in Figure 4, where we plot the same figure of merit, N/N_{mdp}, as

before, but now using $\beta = a_{best}/\sigma_{a,2d}$. This curve is again quadratic in β but rises more steeply than the 1d relation, because $\sigma_{a,2d}$ is larger than $\sigma_{a,1d}$. In agreement with Weisskopf et al. (2010) we find that the 2d relation is very well approximated as $N/N_{mdp} = \beta^2/4.1$, and this is the dashed line running through the square symbols.

We also show in Figure 7 the position angle uncertainty, σ_ϕ in degrees (at 1σ confidence) as a function of $\beta = a_{best}/\sigma_{a,2d}$. We find that this relation can be very well approximated as $\sigma_\phi = 28.5(\mathrm{deg})/\beta$. Thus, a joint 3σ measurement constrains the position angle to better than 10 degrees.

5. Discussion and Summary

The MDP is a very commonly employed figure of merit to describe the polarization sensitivity of a detector system. In their recent paper Weisskopf et al. (2010) have argued that "more counts would be needed" to *measure* the polarization corresponding to the 99% confidence MDA rather than to just establish the same level. As we have shown this is correct in the sense of a "polarization measurement" being a joint measurement of both the amplitude a and position angle ϕ. This seems sensible, since it adds an additional requirement, that the measured a and ϕ both fall within a 2d confidence region. We have also confirmed that the additional number of counts required is given by a factor of ≈ 2.2.

However, if one were to ask a different question, and were interested in establishing simply that a source is polarized, with say, a 3σ measurement of the amplitude, a, without concern for the position angle ϕ, then the 1d (one interesting parameter) case is appropriate, and no extra counts are needed to measure the MDA. So, in some sense the answer one obtains depends on the question asked.

The MDP value has often been used to estimate observing times required to reach

particular sensitivity levels. The results shown here demonstrate that exposure requests should clearly match the measurement goals of the desired scientific program. For example, if source modeling requires a joint measurement of both polarization parameters, then the appropriate time to reach a required precision for two interesting parameters should be requested.

REFERENCES

Bellazzini, R., Angelini, F., Baldini, L., et al. 2003, Proc. SPIE, 4843, 372

Black, J. K., Deines-Jones, P., Hill, J. E., et al. 2010, Proc. SPIE, 7732, 25

Black, J. K., Deines-Jones, P., Jahoda, K., Ready, S. E., & Street, R. A. 2004, Proc. SPIE, 5165, 346

Costa, E., Soffitta, P., Bellazzini, R., et al. 2001, Nature, 411, 662

Dean, A. J., Clark, D. J., Stephen, J. B., et al. 2008, Science, 321, 1183

Forot, M., Laurent, P., Grenier, I. A., Gouiffès, C., & Lebrun, F. 2008, ApJ, 688, L29

Kestenbaum, H. L., Cohen, G. G., Long, K. S., et al. 1976, ApJ, 210, 805

Pacciani, L., Costa, E., Di Persio, G., et al. 2003, Proc. SPIE, 4843, 394

Rice, S. O. 1945, Bell System Technical Journal 24, 46-156

Weisskopf, M. C., Silver, E. H., Kestenbaum, H. L., Long, K. S., & Novick, R. 1978, ApJ, 220, L117

Weisskopf, M. C., Cohen, G. G., Kestenbaum, H. L., et al. 1976, ApJ, 208, L125

Weisskopf, M. C., Elsner, R. F., & O'Dell, S. L. 2010, Proc. SPIE, 7732, 11

This manuscript was prepared with the AAS LaTeX macros v5.2.

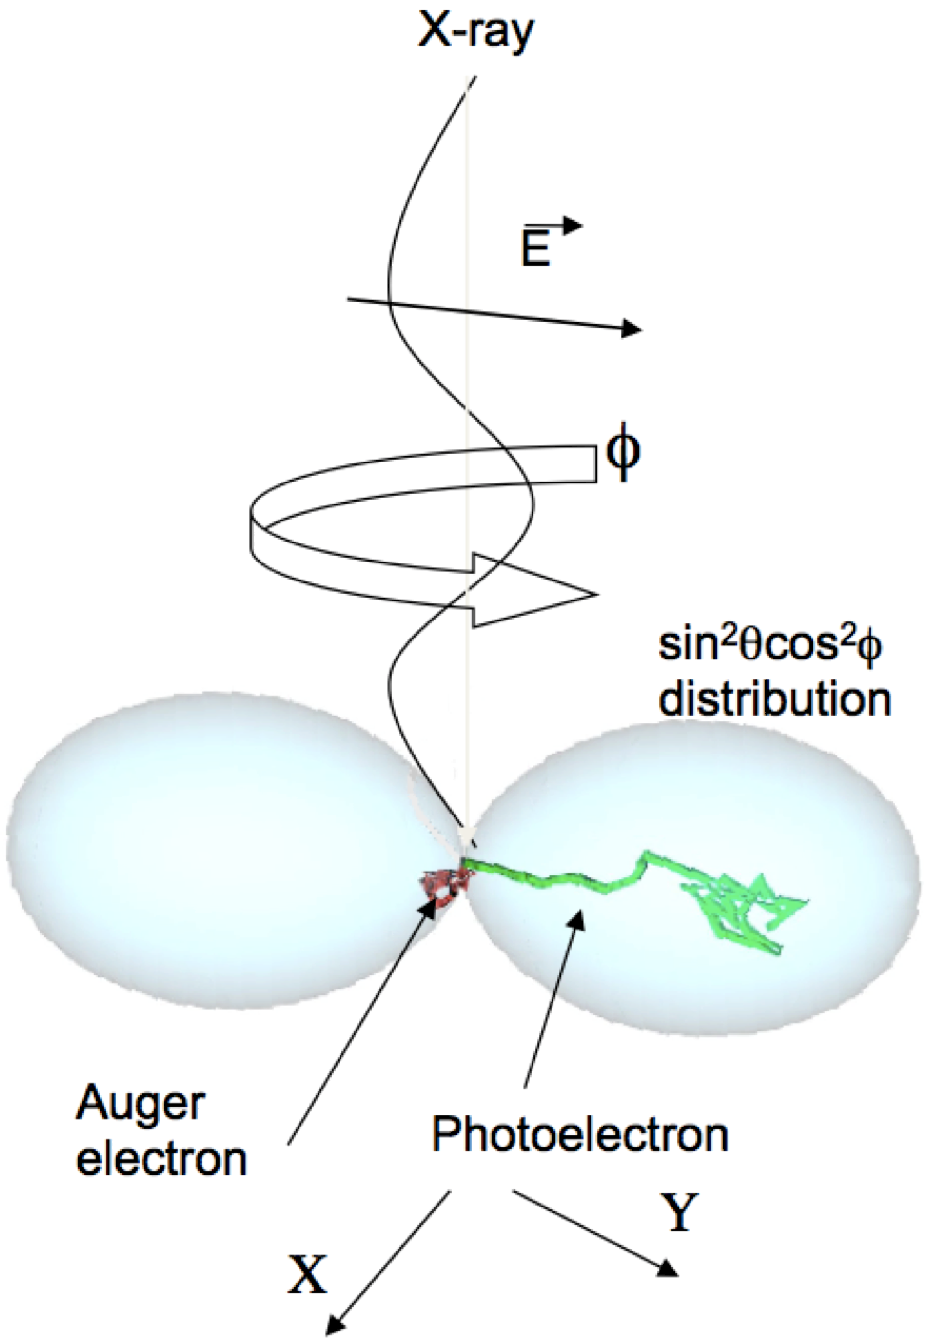

Fig. 1.— Geometry relevant for photoelectric polarimeter measurements. Photons travel from figure top to bottom. Photoelectrons are preferentially emitted in the plane perpendicular to the photon direction of travel (in this case the X - Y plane), and their initial direction is measured by the azimuthal angle ϕ.

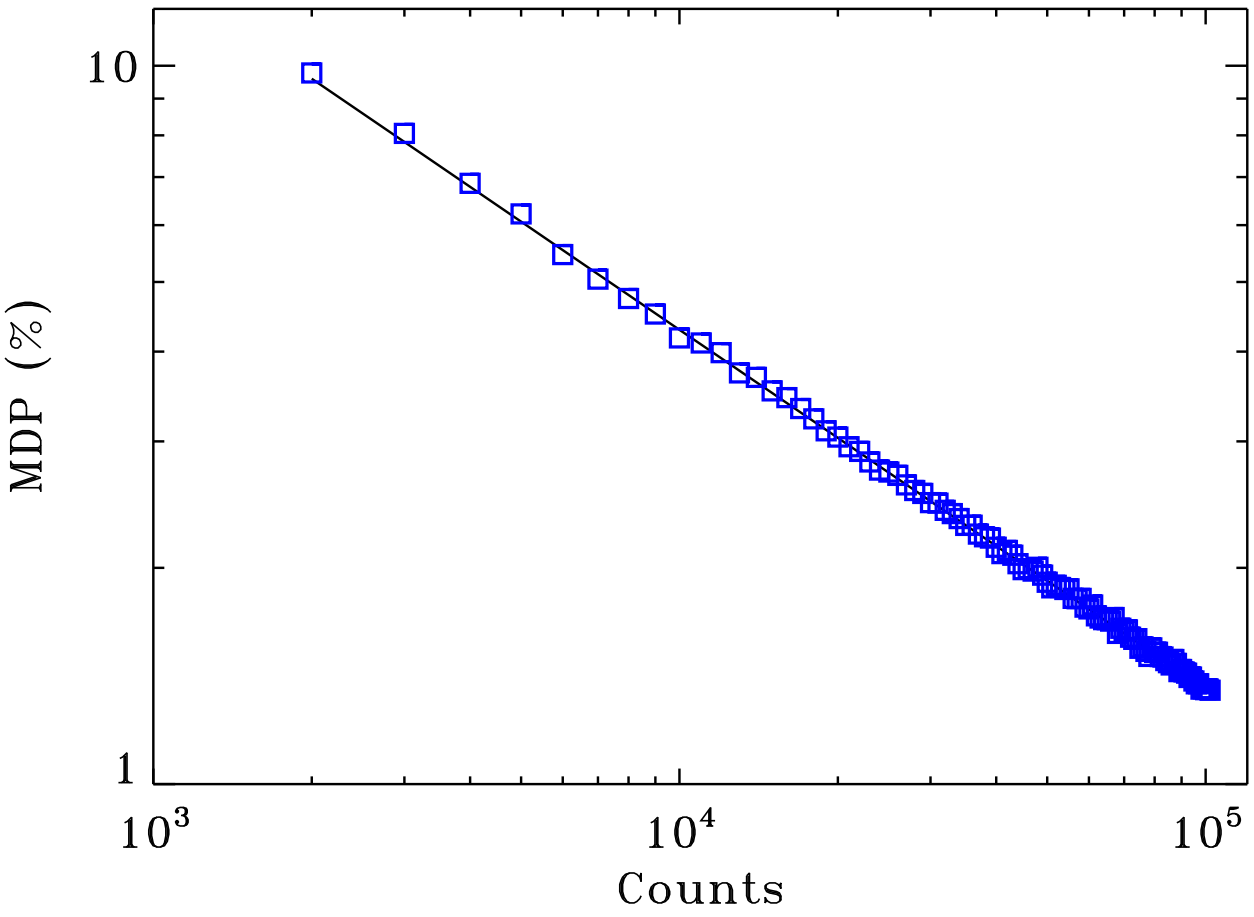

Fig. 2.— The amplitude, $a_{1\%} \equiv MDP(\mu = 1)$, vs the number of counts, N, obtained from simulations described in §4.1 (blue squares), and the analytic formula (solid curve, §3).

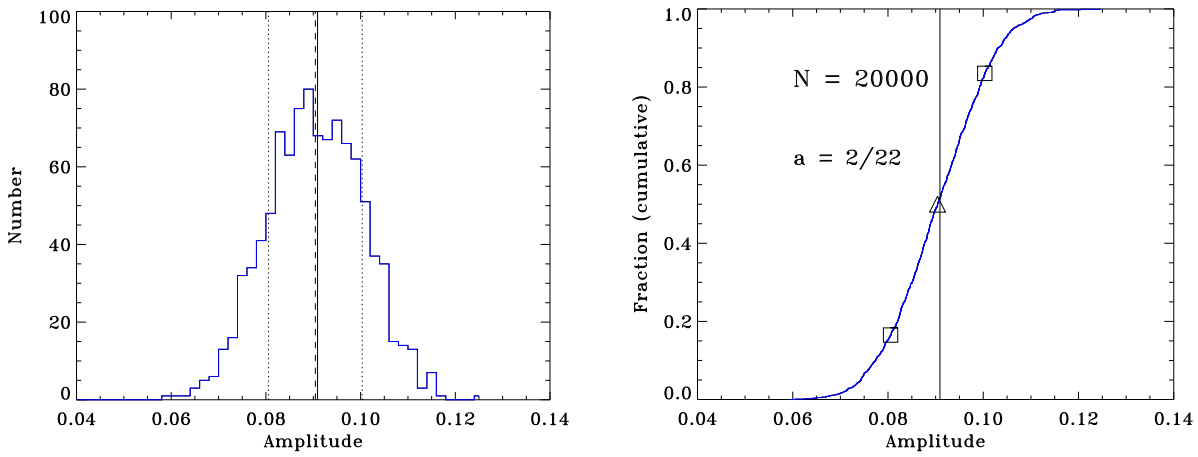

Fig. 3.— Example 1d distributions in the amplitude, a, computed from Monte Carlo simulations. The left panel shows the differential distribution (binned), and the right panel shows the cumulative distribution. To estimate the mean we find the midpoint (triangle symbol), and the 1σ extremes, a_{min} and a_{max} at 0.165 and 0.835, respectively (square symbols, enclosing 68% of the distribution).

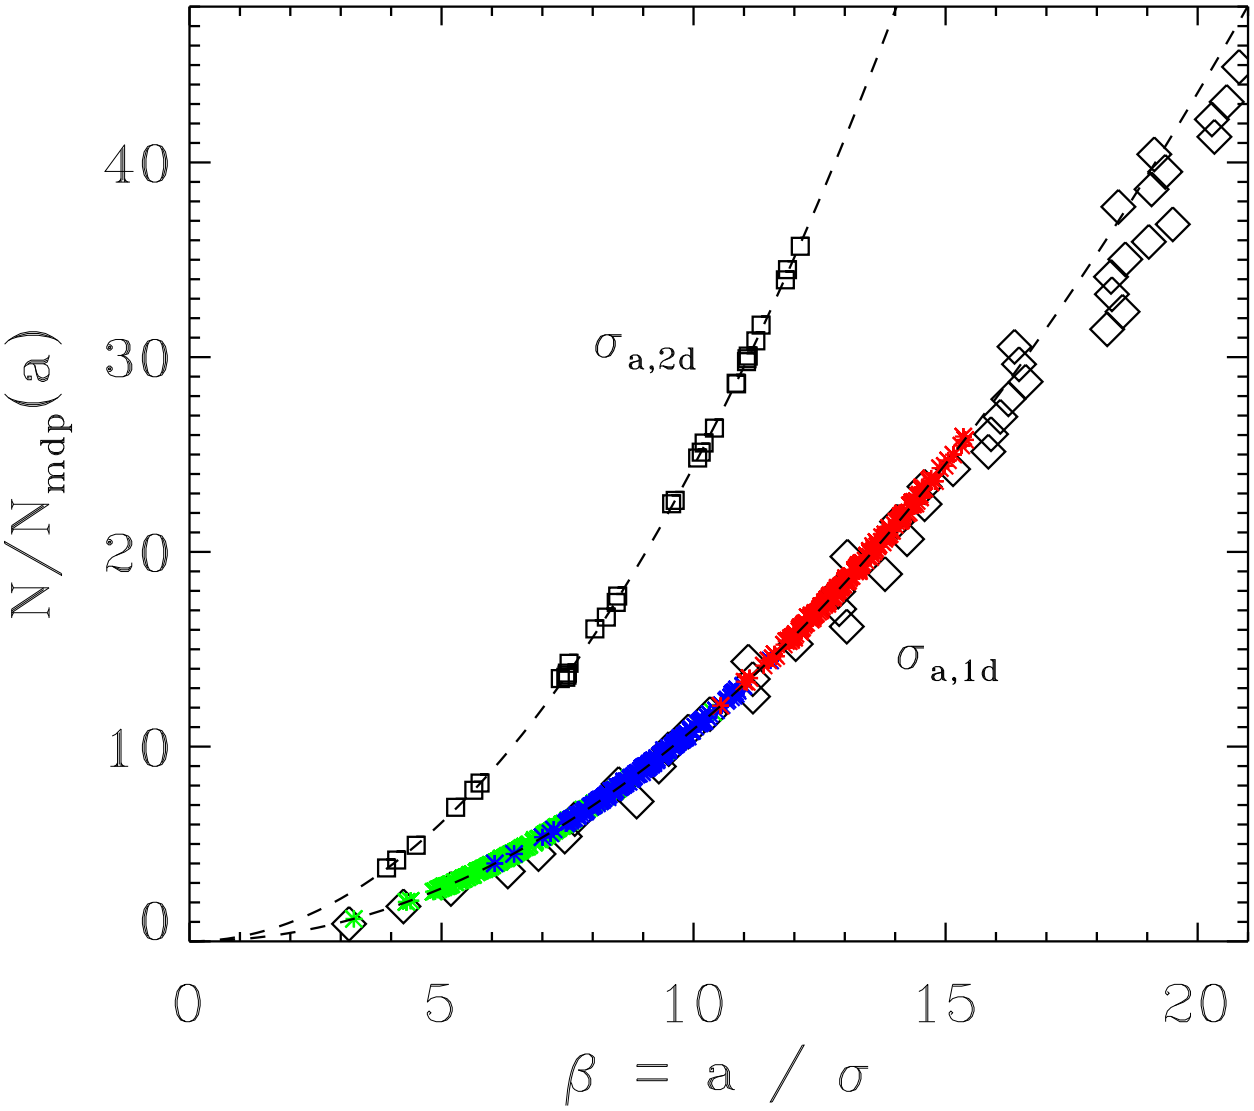

Fig. 4.— Plot of $N/N_{mdp}(a)$ as a function of β, the "number of sigmas" of the measurement. Both the 1d amplitude distribution ($\sigma_{a,1d}$, independent of the position angle) and the 2d joint distribution ($\sigma_{a,2d}$), are shown. The black diamond symbols are derived from the results of $M_{sim} = 1000$ simulations for different values of N (see discussion in §4.3). The colored symbols are the results of individual simulations where β is derived from the best-fit amplitude, a, and its 1d, 1σ uncertainty, $\sigma_{a,1d}$. The red symbols were computed with $N = 10,000$, $a_0 = 2/22$, $\phi_0 = 0$ (deg), the blue used $N = 20,000$, $a_0 = 2/22$, $\phi_0 = 0(deg)$, and the green with $N = 24,000$, $a_0 = 3/25$, and $\phi_0 = 22.5$ (deg). The solid dashed curve for the 1d case is given by $N/N_{mdp}(a) = \beta^2/9.2$. The black square symbols show the 2d confidence region results (see discussion in §4.4). The curve running through the 2d results is given by $N/N_{mdp}(a) = \beta^2/4.1$

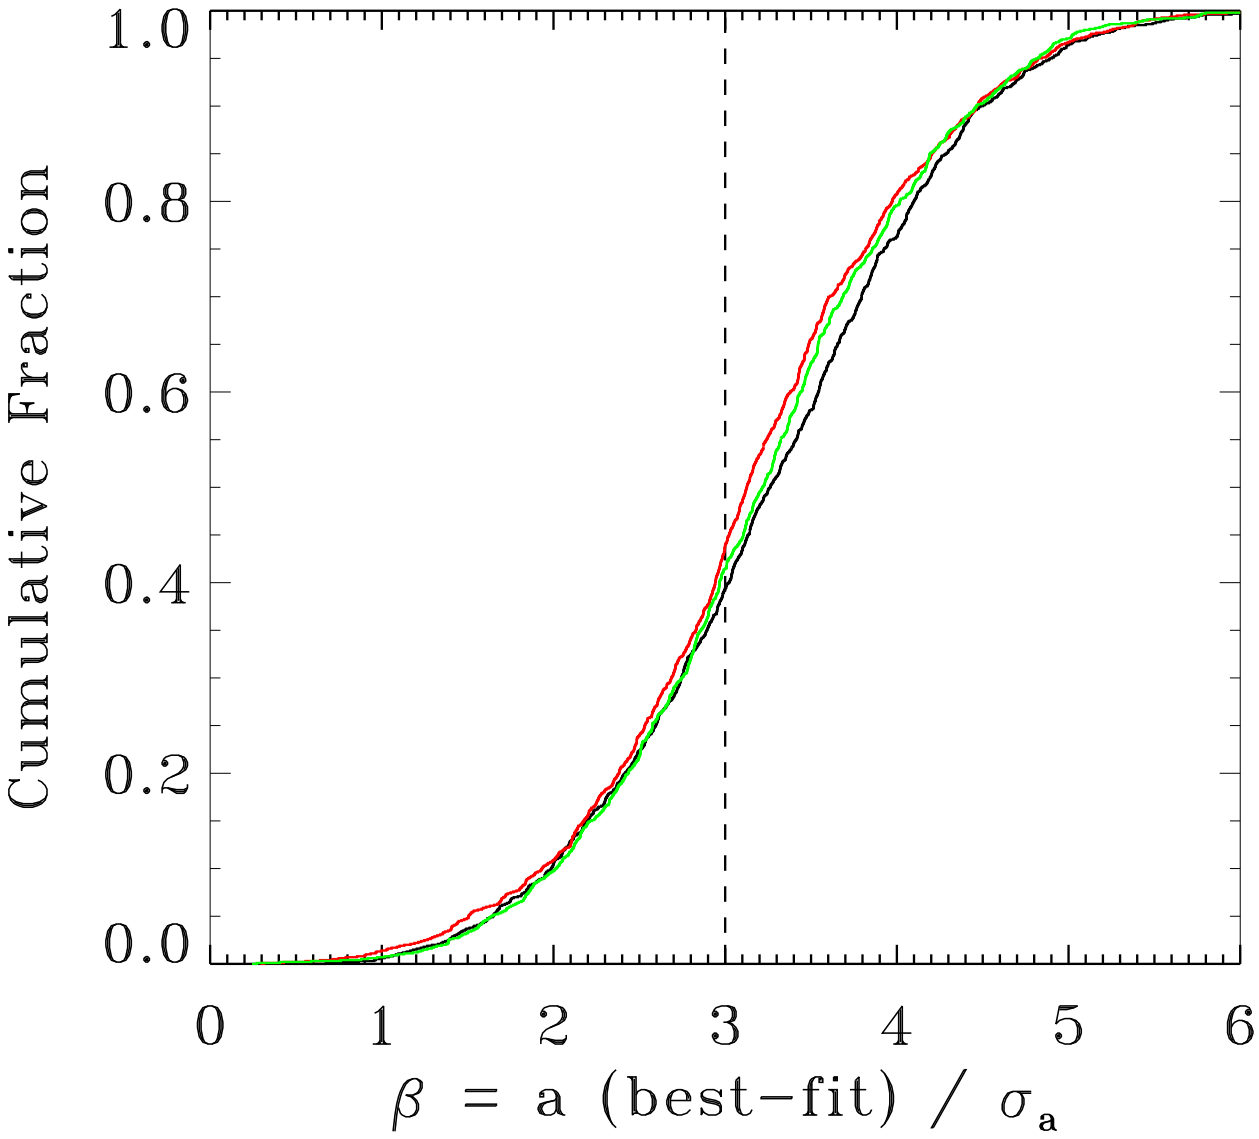

Fig. 5.— Plot of the cumulative distribution of $\beta = a/\sigma_{a,1d}$ for several different simulations all satisfying the condition that a = MDA. A vertical line is plotted at $\beta = 3$ (the nominal 3σ detection criterion), and which is close to the most probable value (the distribution is not exactly gaussian, ie. symmetric). See the discussion in §4.3 for more details.

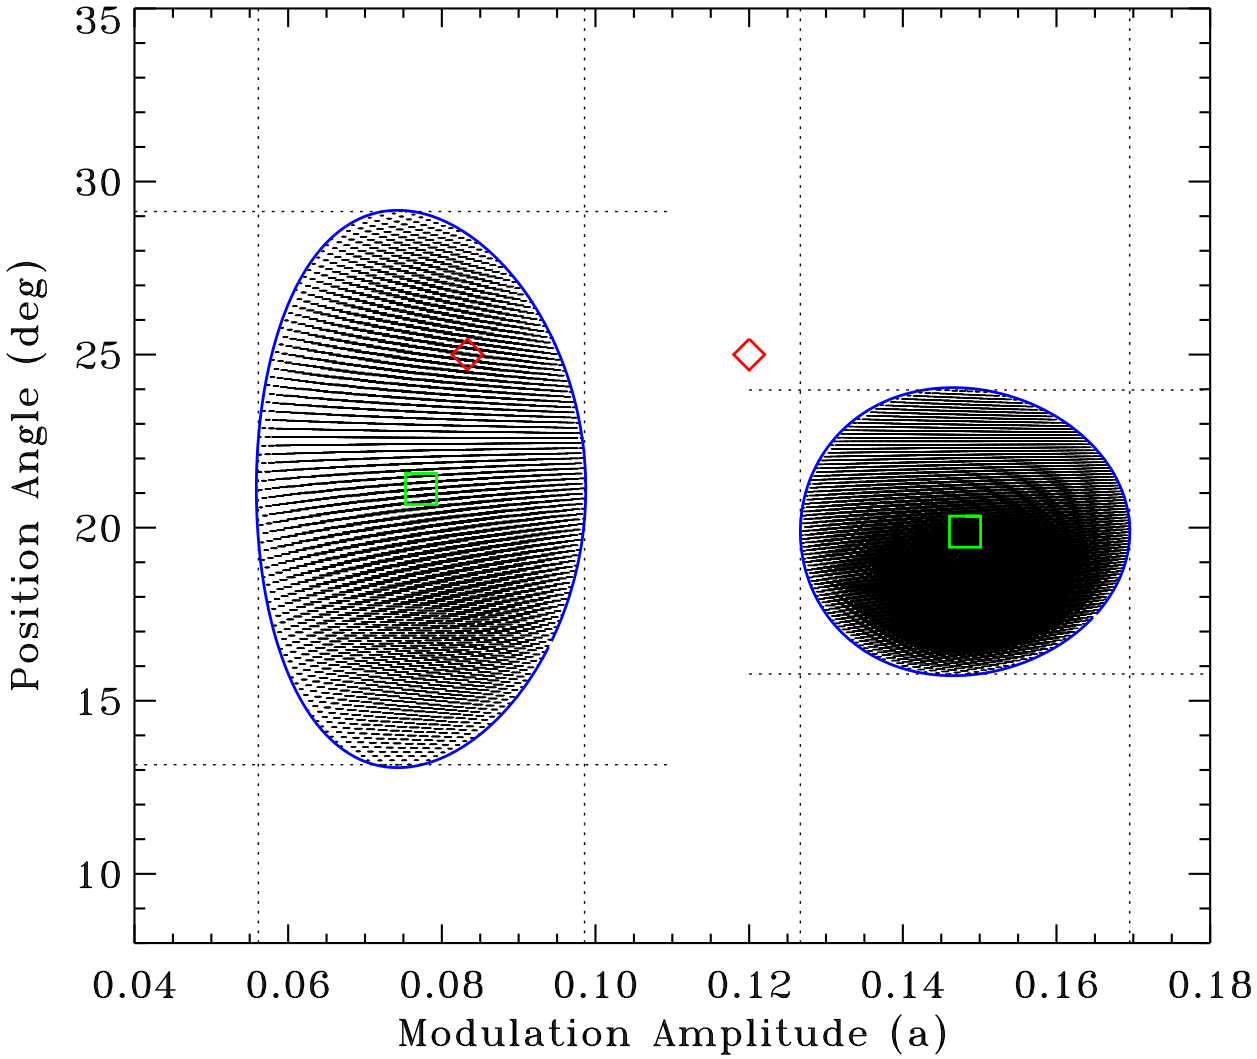

Fig. 6.— Confidence regions and contours in the (a, ϕ) plane. Results from two simulations are shown. Simulated modulation curves were computed with N=10,000 counts for two different intrinsic amplitudes both with a position angle of 25 degrees. These values are marked by the red diamond symbols. The resulting best fit parameter values (green squares) and confidence regions (shaded areas) are shown. The shaded regions are the $\Delta\chi^2 < 2.3$ regions (1σ for 2d confidence regions). The horizontal and vertical dotted lines denote the extremes in each parameter and are used to compute $\sigma_{a,2d}$ and $\sigma_{\phi,2d}$. The blue contour curves were computed from the analytical expressions in §4.4.

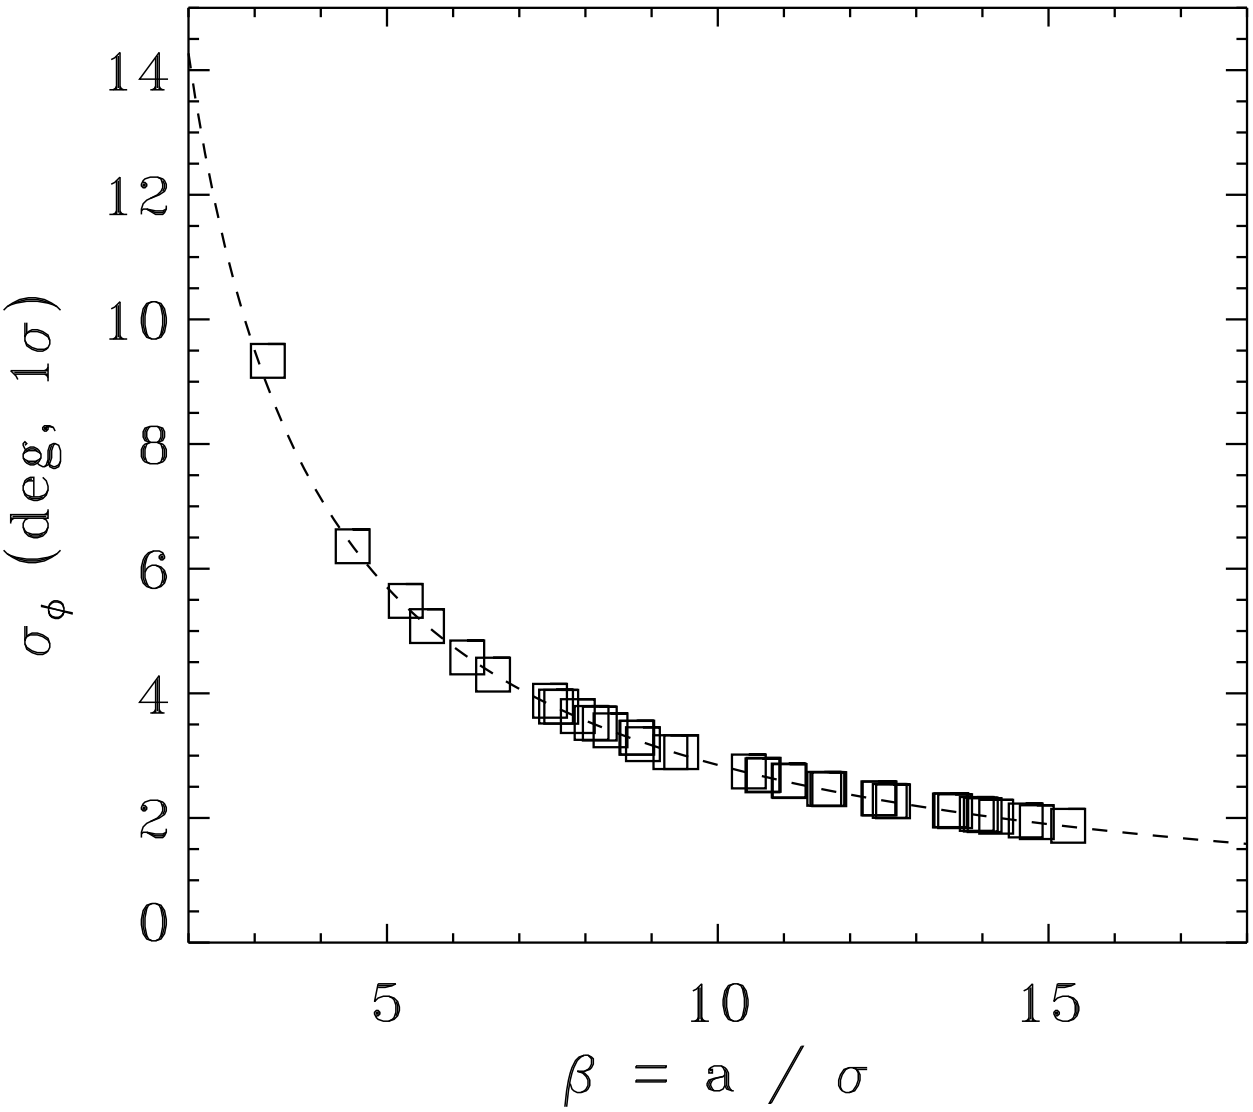

Fig. 7.— Plot of the position angle uncertainty, σ_ϕ (1σ, in degrees), derived from the 2d confidence regions, as a function of β_{2d}, the "number of sigmas" of the measurement. The solid dashed curve is given by $\sigma_\phi = 28.5/\beta_{2d}$.